ARL-TR-8042 ● JUNE 2017

US Army Research Laboratory

Assessment of Energy-Efficient and Model-Based Control

by Craig Lennon, Marshal Childers, Mario Harper, Camilo Ordonez, Nikhil Gupta, James Pace, Ryan Kopinsky, Aneesh Sharma, Emmanuel Collins, and Jonathan Clark

NOTICES

Disclaimers

ARL-TR-8042 ● JUNE 2017

US Army Research Laboratory

Assessment of Energy-Efficient and Model-Based Control

by Craig Lennon and Marshal Childers
Vehicle Technology Directorate, ARL

Mario Harper, Camilo Ordonez, Nikhil Gupta, James Pace, Ryan Kopinsky, Aneesh Sharma, Emmanuel Collins, and Jonathan Clark
Center for Intelligent Systems, Controls and Robotics, Jointly Sponsored by Florida A&M University and Florida State University, Tallahassee, FL

REPORT DOCUMENTATION PAGE	*Form Approved* *OMB No. 0704-0188*

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED (From - To)
June 2017	Technical Report	June 2016

4. TITLE AND SUBTITLE

Assessment of Energy-Efficient and Model-Based Control

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Craig Lennon, Marshal Childers, Mario Harper, Camilo Ordonez, Nikhil Gupta, James Pace, Ryan Kopinsky, Aneesh Sharma, Emmanuel Collins, and Jonathan Clark

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

US Army Research Laboratory
Autonomous Systems Division
Vehicle Technology Directorate (ATTN: RDRL-VTA)
Aberdeen Proving Ground, MD 21005-5069.

8. PERFORMING ORGANIZATION REPORT NUMBER

ARL-TR-8042

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

The US Army Research Laboratory's (ARL's) Robotics Collaborative Technology Alliance is a program intended to change robots from tools that Soldiers use into teammates with which Soldiers can work. One desired ability of such a teammate is the ability to operate in an energy-efficient manner on a variety of surfaces. To develop such a teammate, alliance researchers developed planning algorithms that incorporate knowledge of the vehicle's steering and control system. These algorithms adapt their navigation to different types of terrain, learning appropriate parameter values by conducting a brief set of trial maneuvers, and are intended to enable the robot to operate in a manner that is more energy efficient. In June of 2016, ARL researchers conducted an assessment of this technology by comparing this planning algorithm to a traditional minimum-distance planning algorithm. This assessment found an overall improvement in energy efficiency, which was clearly visible when the systems operated on grass, but unclear when the systems operated on asphalt. Overall, the results suggest that the energy-efficient planner does have the potential to plan more energy-efficient paths.

15. SUBJECT TERMS

unmanned ground system, autonomous systems, robotics, mobility, planning, energy

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
					Craig Lennon
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (Include area code)
Unclassified	Unclassified	Unclassified	UU	30	410-278-9886

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

Contents

List of Figures

List of Tables

1. Introduction

The US Army Research Laboratory's (ARL's) Robotics Collaborative Technology Alliance (RCTA) is an alliance of robotics research institutions working together to transform robots from tools that Soldiers use into teammates with which Soldiers can work. One desired ability of such a teammate is the ability to plan motion over a variety of surfaces in a manner that conserves energy. To further research into this capability, researchers from Florida State University and Carnegie Mellon University developed planning algorithms that select routes based on terrain and the vehicle's motor system. In practice, the robot learns appropriate parameter values by conducting a set of trial maneuvers and then uses these parameters for energy-efficient path planning.

In June of 2016, these capabilities were assessed at Florida State University. The assessment consisted of learning terrain parameters for asphalt and grass, and then navigating a path through obstacles on a flat patch of grass or asphalt. In Section 2, we describe the integrated system and the algorithm that was assessed. Section 3 describes the methodology of the assessment, Section 4 presents the assessment results, and Section 5 contains our conclusions.

2. Robotic System

In this section, we describe the equipment and technology that played a key role in the assessment. We start with the robotic platform, including the sensors and computational power, followed by a description of the planning and learning algorithms.

2.1 Robotic Platform

The robot used in the integrated assessment is a Clearpath Husky, equipped with 3 Mac Mini machines, each with 2.3 GHz quad-core processors, used for hardware control, planning, maintaining a map of the environment, and communicating with the user. A fourth computer with a 1.2 GHz Intel Core 2 DUO processor ran a visual odometry algorithm. The input for the visual odometry was provided by the upper, and downward facing, Bumblebee 2 stereo camera shown in Fig. 1. A Hokuyo lidar was used for obstacle detection. Linear Hall effect sensors were attached to each of the motor's sides to measure wheel speed for the learning phase.

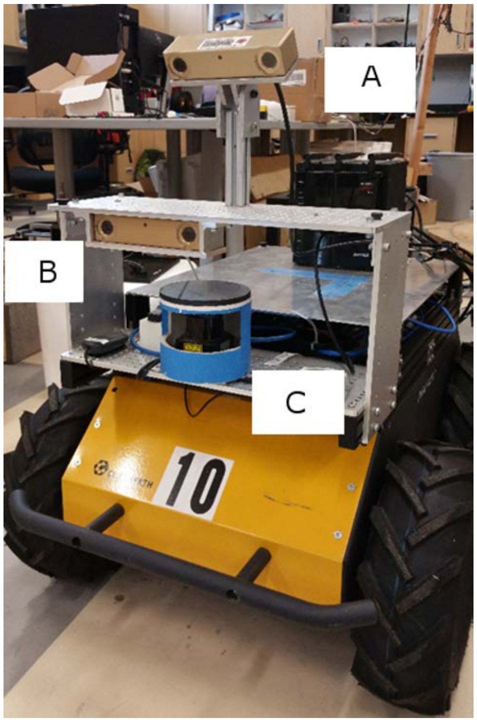

Fig. 1 ClearPath Husky used for the assessment. The upper camera (A) was used for visual odometry, and the lower (B) was not used at all. The lidar (C) was used for obstacle avoidance.

2.2 Dynamically Feasible Energy-Efficient Motion Planning

Skid-steer robots are commonly used as robotic platforms because they are simple and robust. But in order to turn, the wheels of such a vehicle must skid or slip, a terrain-dependent action that complicates localization.[1] Sharp turns may also require motor torque in excess of what the vehicle can provide, which can lead to route re-planning and wasted energy.[1] Gupta et al.[2] proposed a method of planning that constrains paths to those that are dynamically feasible given a terrain type, payload, and vehicle model. It does so by estimating the required torque for a turn given the terrain, and then sampling possible path extensions, rejecting those that demand torque in excess of what the vehicle could provide. This search is guided by a heuristic with a preference for energy-efficient paths. To estimate the torque required for a turn, the model requires the setting of parameters related to the platform structure and mass, and the resistive forces of friction. That means that given a known, skid-steered platform and known payload, the model requires only parameters depending on terrain type.

2.3 Learning of Terrain-Dependent Models

The robot begins with a kinematic and a friction model that has been learned based on a theoretical model and prior navigation. When the robot faces new terrain, the system drives along a diagnostic path, varying speed and turning radii. The robot then updates its previously learned models based on its more recent navigation experience. The physical experiments to date have been conducted on flat terrain. The theoretical model in Gupta et al.[2] was constructed to support navigation on a plane with constant slope, although physical experiments have been conducted only on terrain with zero or negligible slope. For terrain with varying slope, the terrain could likely be approximated by the robot as a union of segments of planes, each with constant slope, and with the sampling algorithm choosing sample paths within plane segments. Alternatively, an approximate plane could be fit to the terrain, the path calculated on the approximate path, and then executed on the actual terrain. In this case, areas of actual terrain with the greatest difference from the approximated plane could be avoided by using the residuals of the fit as part of the planning heuristic.

2.4 Purpose of the Experiment

These algorithms have been integrated with the RCTA world model[3] that is examined in Lennon et al.,[4] so that learning of terrain models and energy-efficient motion planning can be performed online by the RCTA system. This experiment was intended to compare minimum distance motion planning with the energy-efficient planning method from Ordonez et al.[1] to determine if the latter was more energy efficient. A secondary purpose was to verify that the learning process of Section 2.3 does improve the planning described in Section 2.2.

3. Experimental Design

The design of the experiment was effected substantially by environmental conditions. A sudden tropical storm brought rain during the week of the assessment, with the result that the assessment needed to be compressed into 1 day to avoid the rain. This limited the number of runs we could conduct, but did not change the basic structure of the experiment. The assessment consisted of navigation over a flat patch of asphalt or grass, with a series of obstacles placed on it. The key factors and variables are listed in Table 1. If more time had been available, more terrain types would have been included. The terrain type was not originally thought to be a factor related to planner performance, but only a source of variation in performance, along with obstacle configuration. Obstacle configurations were not kept the same between the grass and asphalt configurations, because this would have required more time to carefully place obstacles, and we believed that this time was better spent by getting in more runs.

Table 1 Variables and factors

Factor/Variable	Levels
Planning type	Minimum distance, energy efficient, energy efficient without learning
Terrain	Asphalt, grass
Obstacles	Configurations of obstacles were treated as blocks
Energy used	The energy in joules used during the run (continuous variable)
Collisions	Count of the number of collisions with obstacles (whole number)

The terrain types considered were asphalt and grass, and the obstacles were cardboard cylinders that the robot would need to steer around. Each separate configuration of cylinders was treated as a block, with the intent of examining the difference in performance of the algorithms within each block. Since the primary purpose was to compare energy-efficient planning with learning against minimum-distance planning, the planning type was kept at 1 of these 2 levels for 32 out of the 36 runs. The overall number of runs (36) was determined by the time restrictions. An example of a set up course is shown in Fig. 2, and the design is shown in Table 2. In this table, configuration refers to a given setup of barrels, and the algorithm values are: minimum distance (MD), energy efficient (EE), and energy efficient without learning parameter values (EE Default). The intent was to allow a comparison between the amount of energy used by the minimum-distance and energy-efficient planners within each block. This would provide 16 comparison values to assess the difference between the minimum-distance and energy-efficient planners, and 4 comparison values for energy-efficient planning with and without learning. The order of the algorithms was randomized within the blocks. We did not originally intend to compare the performance of the planners by terrain type.

Fig. 2 Representative course layout. The goal position is marked by an orange cone.

Table 2 Design for the experiment

Run	Terrain	Configuration	Planning algorithm[a]
1	Asphalt	1	MD
2	Asphalt	1	EEP
3	Asphalt	2	EEP
4	Asphalt	2	MD
5	Asphalt	3	EEP
6	Asphalt	3	MD
7	Asphalt	4	EEP
8	Asphalt	4	MD
9	Asphalt	5	EEP
10	Asphalt	5	MD
11	Asphalt	6	MD
12	Asphalt	6	EEP
13	Asphalt	7	MD
14	Asphalt	7	EE Default
15	Asphalt	7	EEP
16	Asphalt	8	EE Default
17	Asphalt	8	MD
18	Asphalt	8	EEP
		Change locations	
Run	**Terrain**	**Configuration**	**Algorithm**
19	Grass	9	EEP
20	Grass	9	MD
21	Grass	10	MD
22	Grass	10	EEP
23	Grass	11	MD
24	Grass	11	EEP
25	Grass	12	MD
26	Grass	12	EEP
27	Grass	13	MD
28	Grass	13	EEP
29	Grass	14	EEP
30	Grass	14	MD
31	Grass	15	EEP
32	Grass	15	EE Default
33	Grass	15	MD
34	Grass	16	EE Default
35	Grass	16	MD
36	Grass	16	EEP

[a] Algorithm was minimum distance (MD), energy efficient (EE), or energy efficient without terrain learning (EE Default).

The most important dependent variable was the actual energy expended during the run, but the following were also recorded:

- whether a collision occurred
- time taken to finish the run (Time)
- number of times re-planning occurred to avoid collision (CollisionReplan)
- number of times tracking re-planning occurred. These re-plans are generally caused by localization error because of issues with perception, slipping, or skidding (TrackingReplan)

The Appendix includes recorded data from all runs (Table A-1) and all variables and their definitions are listed in Table A-2.

The expectation was that with energy-efficient planning, CollisionReplan and TrackingReplan should be no greater than the number of such occurrences under minimum-distance planning, and that time should not differ substantially between the 2 planning methods. These comparisons were not central to the experiment, and we planned to subject them only to exploratory data analysis, without any direct testing.

4. Results

The experiment occurred on June 7, 2016, during a break in a tropical storm, with the runs on the asphalt lasting from about 0900 to 1200, and the runs on the grass lasting from about 1300 to 1600. We break our examination of results into subsections, with the first being a comparison of the 2 planners with respect to the difference in energy expenditure and the difference in the number of collisions with obstacles. The second is an exploration of the time required to complete the course and the number of re-planning attempts.

4.1 Energy and Collision

The experiment was designed to compare minimum-distance and energy-efficient planning on the basis of the energy expended (in joules). For energy expended, we intended to look at each configuration and consider the difference in energy usage between the 2 planning methods. We had originally considered using a paired t-test to examine the significance of the differences; however, we found that the variability of performance of the planners, possibly due to perception, made the power of such a test low. The results of Gupta et al.[2] suggested an energy savings of 30%–40% might be possible on inclined planes, and the expectation was that an energy savings of no more than that would be possible on a flat surface with a real system. Given such a savings, a paired t-test with 16 pairs, and with the variance actually observed, would have had a power of 0.99, thus failing to appropriately reject a null hypothesis of equality with probability 0.01 (at a 0.05 level of significance). However, if the energy saving was only 10%, the power would drop to 0.8, meaning we would fail to detect a true difference in the means with a probability of 0.2. Some of the variability in performance could be attributable to the performance of the algorithms with different terrain configurations. We cannot neglect the possibility, however, that some performance differences are due to perception errors—particularly in the extreme cases we discuss in Section 4.3. This added variability in performance would make it more difficult to detect a true difference in energy savings and would lower the power of the test. Thus, while we included the results of the originally planned paired t-tests, we are inclined to give more weight to the confidence intervals, and in particular to the confidence intervals with the extreme values removed. The difference in energy use between the 2 planners is listed in Table 3 and the results of the previously mentioned tests are in Table 4.

Table 3 Differences in energy expenditure by configuration

Configuration	Difference (MD–EE)	Terrain
1	88.34	Asphalt
2	–876.61	Asphalt
3	–974.22	Asphalt
4	26.59	Asphalt
5	2,919.41	Asphalt
6	–166.71	Asphalt
7	138.96	Asphalt
8	1,453.43	Asphalt
9	–971.19	Grass
10	–828.91	Grass
11	–2,775.33	Grass
12	–922.88	Grass
13	–2,819.5	Grass
14	–5,394.35	Grass
15	1,033.75	Grass
16	–344.33	Grass

Table 4 Results of t-tests of actual energy use

	All observations	Extremes removed
Number of pairs	16	14
p-value	0.13	0.06
Confidence interval for EE planner energy savings	(–262, 1812) joules	(–39, 1458) joules
Average savings for the energy-efficient planner	775 joules	710 joules

We also found that we needed to consider the possibility of a difference by terrain type. During exploratory data analysis, we noticed a difference in planner performance on grass and asphalt. The density of the differences in energy expenditure are shown in Fig. 3, separately for grass and asphalt. The amount of energy expended by each planner on each run is shown in Fig. 4. As shown in Figs. 3 and 4, the performance of the energy efficient planning seemed to be superior on grass, but no difference is discernable on asphalt. There was also a noticeable difference between the overall amount of energy expended on grass and on asphalt, with runs on asphalt using more energy than those on grass, as shown in Fig. 5.

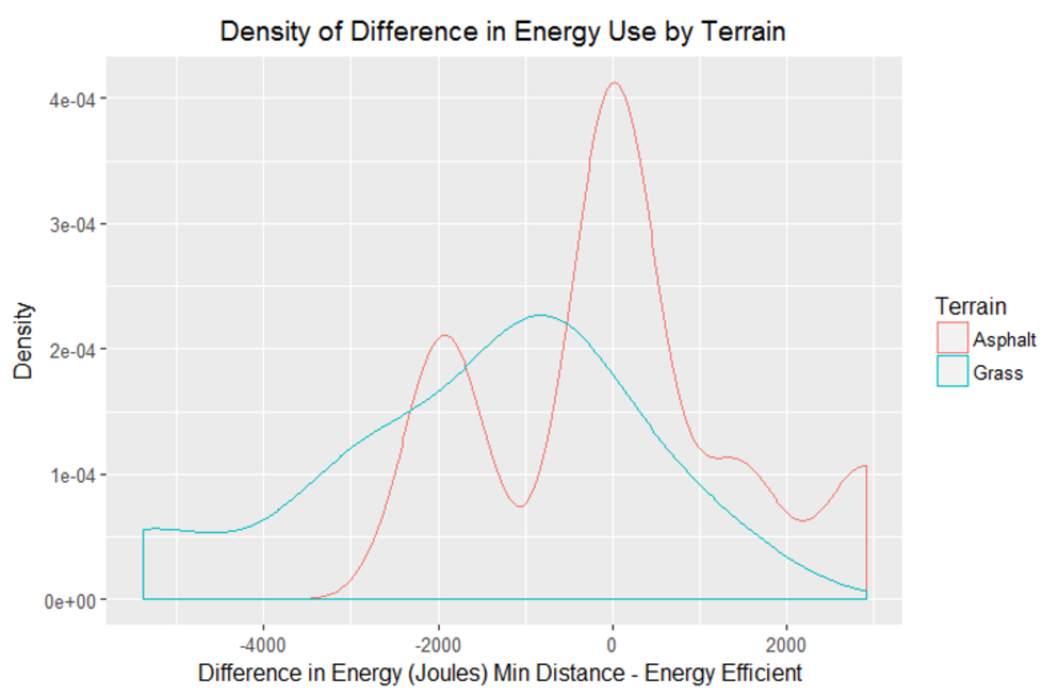

Fig. 3 Density of difference in energy expenditure within configurations by terrain

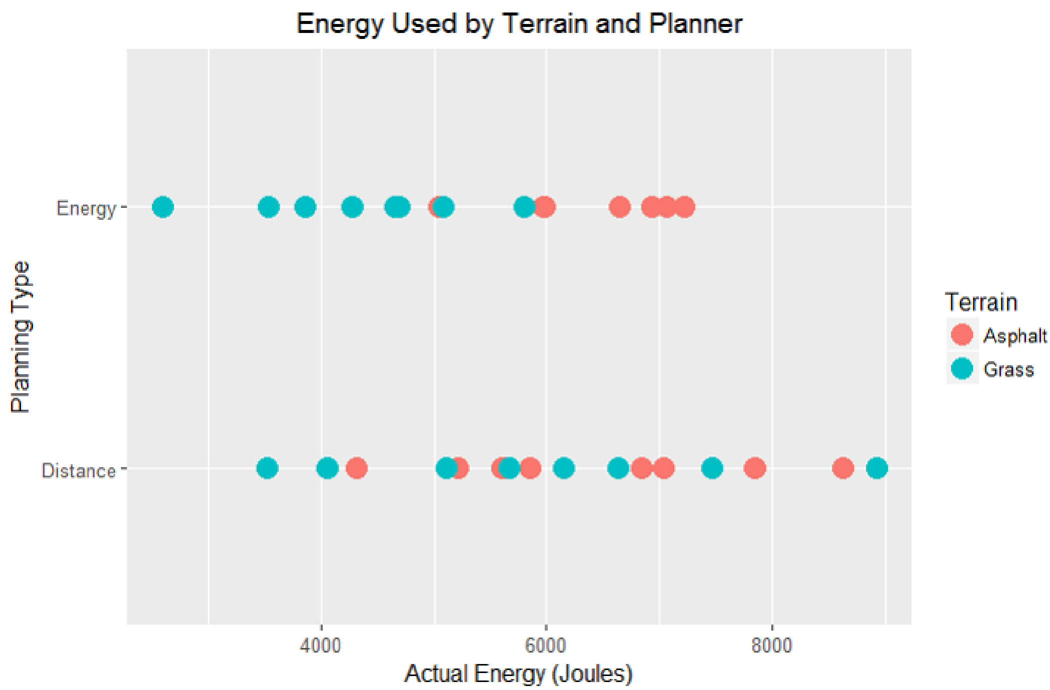

Fig. 4 **Energy expenditure by terrain and planning type**

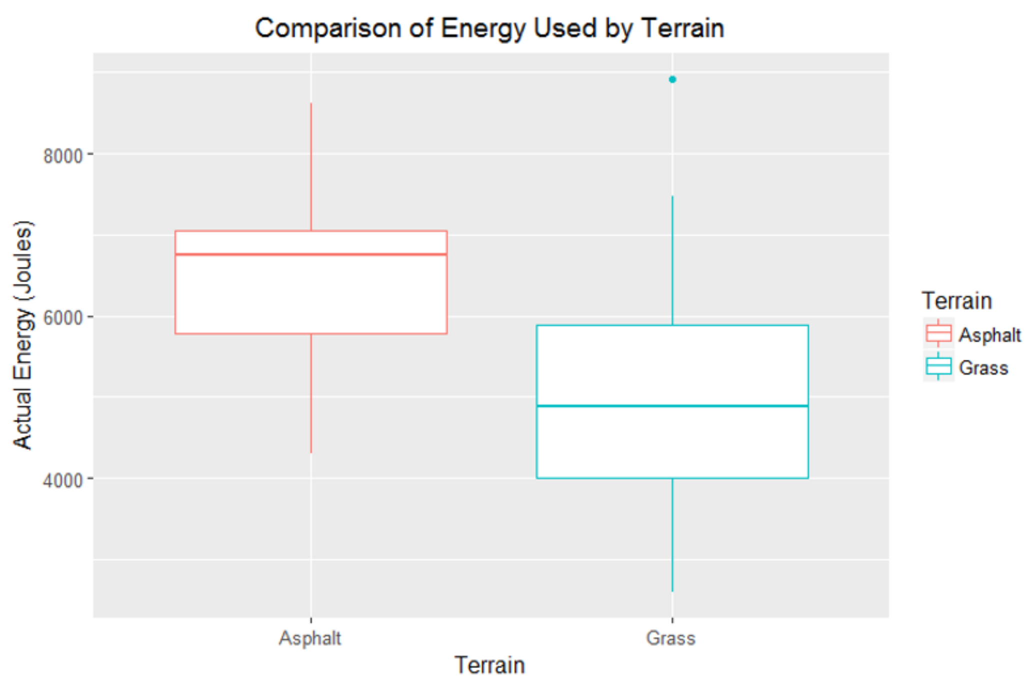

Fig. 5 **Boxplot of aggregate energy use by terrain**

We conducted the planned 2-sided t-test with a null hypothesis of no difference in the mean of actual energy expenditure between the 2 planners. This test neglects terrain and considers all differences as a sample from a single distribution. From this test, we obtained a p-value of 0.1319 and a 95% confidence interval of –262, 1812 joules, which represents a confidence interval for the energy that the energy-efficient planner would save, on average. In the process of checking our model, however, we examined a quantile–quantile (Q–Q) comparison with a normal distribution, shown in Fig. 8 of Section 4.3. This brought to our attention 2 extreme observations that while not in error, deviated substantially from the others. One represented the best performance of the energy-efficient planner relative to the minimum-distance planner, and one represented the worst performance by the energy-efficient planner relative to the minimum-distance planner. We conducted the t-test again, with these samples removed, and found that with this modification, one would find evidence of a difference between the planners, with a p-value of 0.06 and a 95% confidence interval from –39, 1458 joules, again representing a confidence interval for the energy that would be saved by using the energy-efficient planner, on average.

Based on the exploratory data analysis (Fig. 3), and on an examination of the collision data, which we will present later in this subsection, we also conducted an analysis of variance (ANOVA) to consider the interaction of terrain and planner. Because of the pairing, we conducted a 2-factor repeated measures ANOVA, thinking of each configuration as a subject and each planner as a treatment applied to each subject. Thus, the planner is considered a within-subject effect and the terrain was considered a between-subjects effect. The results are shown in Table 5 and provide weak evidence for an interaction of terrain and planner type. While this could be due to the actual terrain, it could also represent the inadequate learning of the parameters for that terrain type, since parameter learning was conducted only at the start of the runs on a given terrain type.

Table 5 ANOVA results

	DF	Sum square	Mean square	P value
Terrain	1	1,999,605	1,999,605	0.2950
Planner	1	64,700	64,700	0.8490
Terrain: Planner	1	5,132,349	5,132,349	0.0988
Residuals	26	45,516,174	1,750,622	. . .

With the extremes of performance of the energy-efficient planner removed, there is evidence of an energy saving advantage to energy-efficient planning. Even with the extremes, however, the confidence intervals of Table 4, suggest that, if energy-efficient planning and minimum-distance planning use equal resources to perform, one might expect, on average, to achieve energy savings from the energy-efficient planner. These confidence intervals are more informative than the p-value of the test because over a course of such short length there is unlikely to be a great difference in energy expenditure, making a difference between the planners difficult to detect, especially with variability added by other factors. Indeed, for the probability of falsely rejecting the null hypothesis to be equal to the probability of incorrectly failing to reject the null hypothesis (power of 0.9 at a significance level of 0.1), the energy-efficient planner would have needed a mean energy savings of 1210 joules (20% over) the minimum-distance planner, which is likely difficult to obtain over a course with a length of 10 m.

In path planning, obstacle avoidance can be achieved by placing a buffer around detected obstacles, and prohibiting the planner from planning a path within the buffered area. The smaller the buffer, the greater the likelihood of collision, and the larger the buffer, the less space available for planning paths. In this experiment, this buffer was set to be small, allowing for more flexibility in planning paths, but greater likelihood of collision. Therefore, although several collisions occurred with the minimum-distance planner, such collisions are not inherent to minimum-distance planning, and could be avoided with greater constraints on the planner, which would presumably result in longer paths and higher energy use. However, since both planners used the same buffer around objects, we considered it reasonable to compare them based on these criteria.

When, during the course of a run the robot touched an obstacle, the robot's starting coordinates were checked to make sure that no error was caused by drift over the course of the experiment. If the starting coordinates were incorrect, the robot was rebooted and the run repeated. If the coordinates were accurate, a collision was recorded and the run counted. Table 6 shows the number of collisions by planner and terrain, and these results suggest an interaction of terrain and planner. Table 7 shows the collision totals by planner type. While we cannot be certain of the reasons behind the greater number of collisions on the grass, the results of Table 6 and Figs. 4 and 5 suggest a plausible reason: the grass was slippery and the asphalt was not. Slippery grass would make collisions more likely (Table 6) and would make turns less energy intensive, leading to less energy use on the grass (Fig. 4). Moreover, this slipperiness would be taken into account by the energy-efficient planner, after its learning phase, and not by the minimum-distance planner.

Table 6 Collisions by planner and terrain, with x/y indicating collisions in x out of y runs

Terrain/planning method	Min distance	Energy efficient
Grass	5/8	0/8
Asphalt	0/8	0/8

Table 7 Collision results matched for each configuration

		Energy efficient	
		Collision	**No collision**
Minimum distance	Collision	0	5
	No collision	0	11

We consider the combination of Tables 4 and 5 as evidence of a potential advantage of the energy-efficient planner on the grass, and possibly on other slippery terrain. The advantage is not in avoiding collisions—which could be done by buffering around obstacles—but in planning paths that avoid them without such buffering, which represents a possible advantage because such buffering sometimes has the effect of making traversable paths appear untraversable.

As a secondary goal, we wanted to verify that the learning process improved the performance of the energy-efficient planner. Considering the 4 configurations in which this comparison was made, a paired t-test of the actual energy expended provides evidence for rejecting the hypothesis that the mean energy expended is equal in favor of the alternative hypothesis that the difference in the means is not equal to 0 (p-value 0.06), with a 95% confidence interval of (–137, 2664) for the mean-energy savings gained by learning before using the energy-efficient algorithm. The actual values for these 4 runs are in Table 8.

Table 8 Differences in actual energy by energy-efficient planning with and without learning

Configuration	Without learning–With learning (joules)
7	–914
8	–2492
15	–457
16	–1153

4.2 Re-planning and Time

Having made the comparisons for which the experiment was conducted, we briefly explore some secondary concerns: differences in re-planning and in time taken to complete the course. Figures 6 and 7 are boxplots comparing the distributions of the number of re-plans due to trajectory corrections (tracking) and for the purpose of collision avoidance (collision). These plots show that the average number of tracking re-plans was similar, but that the distance planner needed a greater number of collision avoidance re-plans. Fewer collision avoidance re-plans suggests a more reasonable choice of path, but we note that a gain in time from fewer re-plans may be balanced by longer time taken to make the plan. No data were collected to examine this because the tradeoff between frequent-but-hasty planning versus better-but-lengthier planning was not an important aspect of the experiment. Table 9 provides the average number of re-plans by planner, type of re-plan, and terrain, as well as the average length of time to complete the course by terrain and planner type. This table, along with a plot of the completion times in Fig. 8, suggest that the route chosen by energy-efficient planning does not necessarily take longer to complete than that chosen by minimum-distance planning; although the added time due to collisions and the short courses involved make us wary of drawing much from the completion time data. The re-planning and time data provide no evidence of worse performance by the energy-efficient planner.

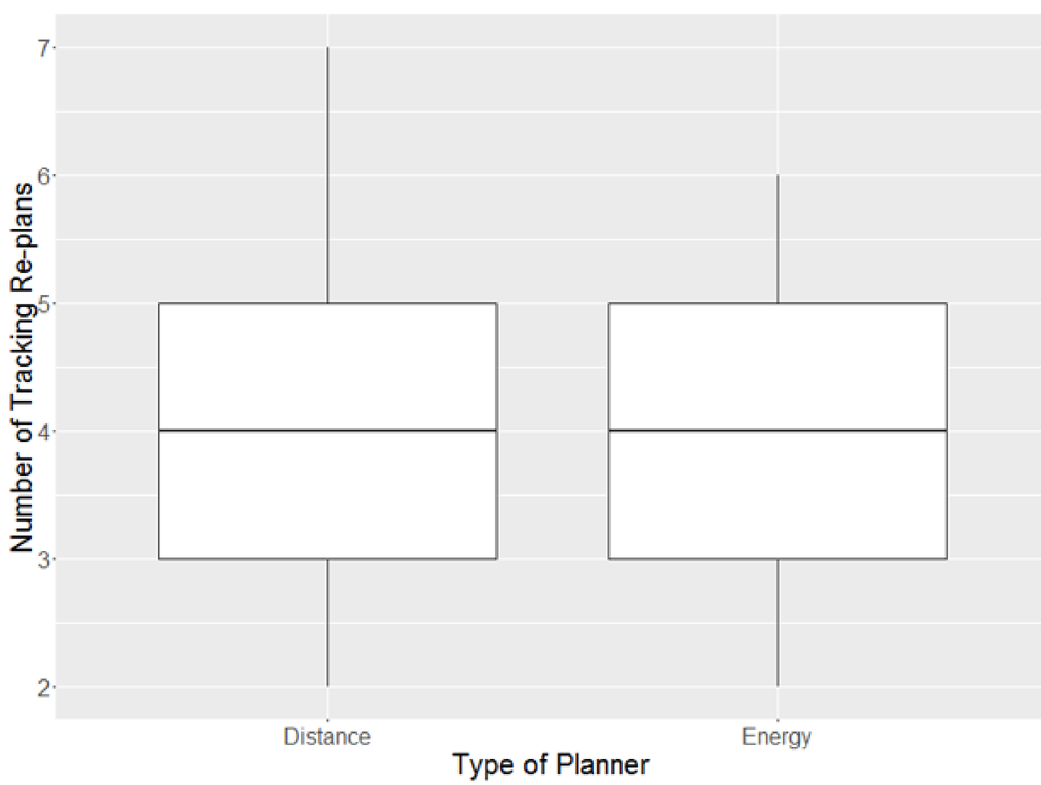

Fig. 6 Number of tracking re-plans by planning type

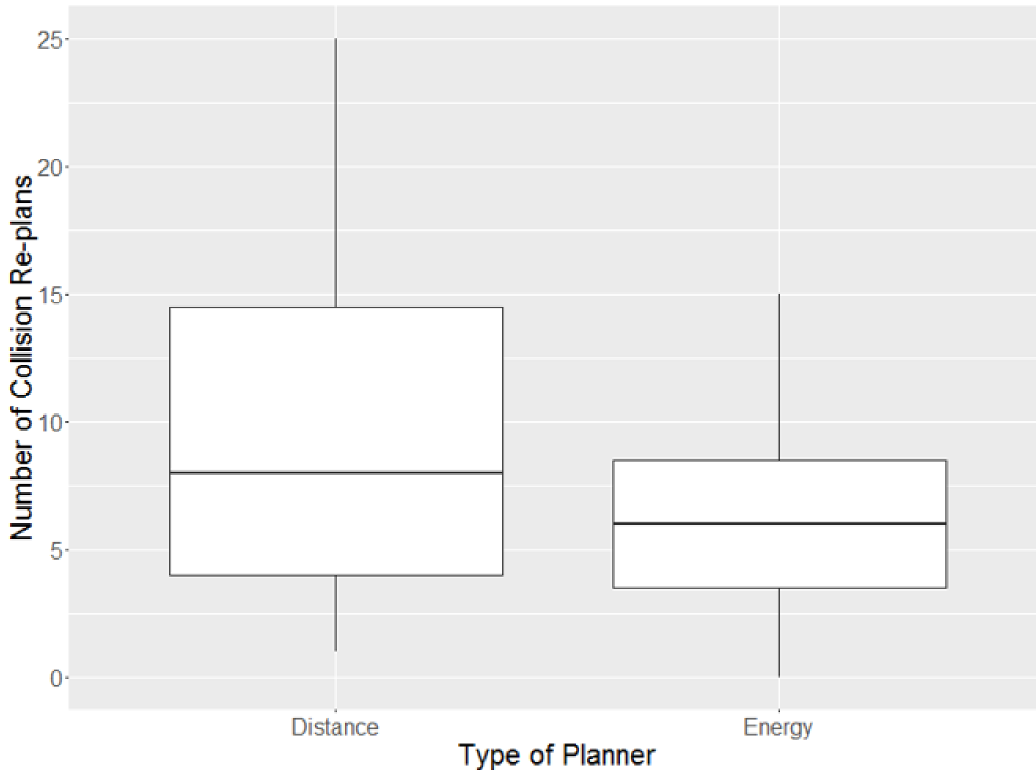

Fig. 7 Number of collision re-plans by planning type

Table 9 Re-plans by terrain and planner

Terrain	Planner	Average no. tracking re-plans	Average no. collision re-plans	Time (s)
Grass	Distance	5.1	9.0	25.1
	Energy	4.0	6.7	24.0
Asphalt	Distance	3.6	10.5	24.9
	Energy	4.1	6.3	24.1

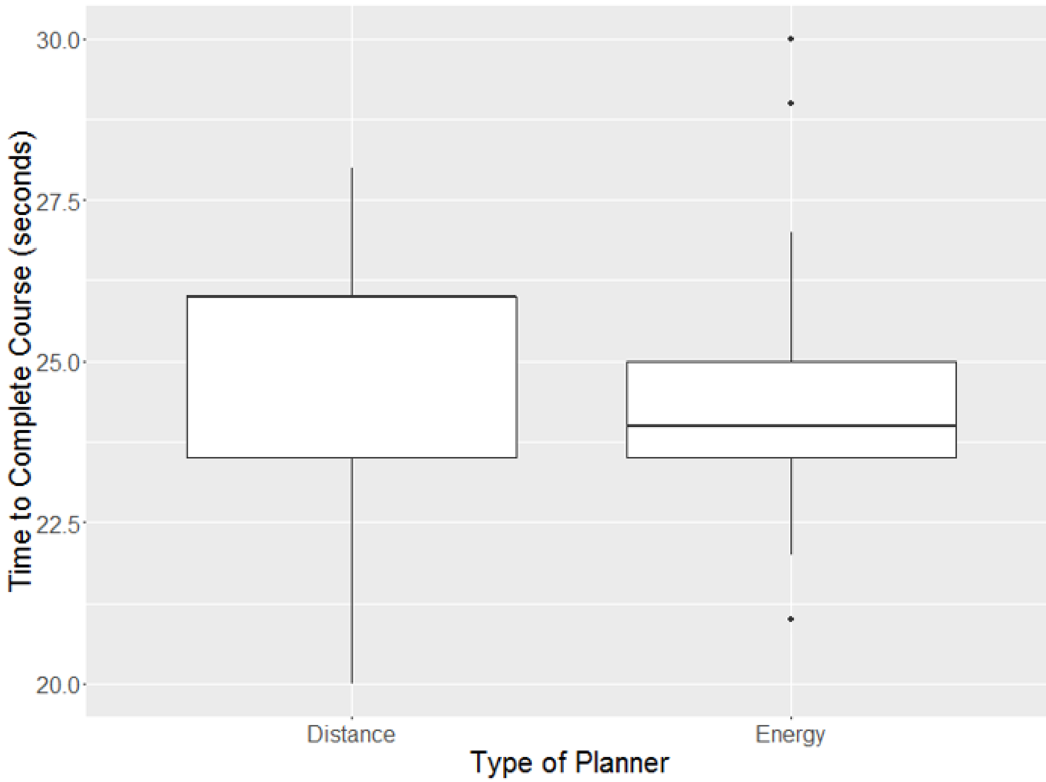

Fig. 8 Time to complete the course by planner type

4.3 Outliers and Model Checking

In examining Table 4, the results from configuration 5 stood out from the other asphalt runs, and the results of configuration 14 stood out from the other grass runs. Consequently, we re-examined the video from these runs to see if anything unusual effected the results. We found nothing unusual in the runs. In configuration 5, the energy-efficient planner just chose a longer route to avoid sharper turns, and paid a price in energy consumption for it. Likewise, in configuration 14 the choice of the energy-efficient planner simply paid off better, both in energy saved and in avoiding barrel collisions. In either case, the odd choice of path could have been influenced by perceptual errors or by errors in other parts of the system.

Our use of the paired t-test for analysis assumes that the differences in energy expended by the minimum-distance and energy-efficient planners are approximately normally distributed. Figure 9 shows a normal Q–Q plot of these differences. The assumption of normality seems appropriate, although the 2 configurations mentioned previously deviate significantly from the line representing a normal distribution.

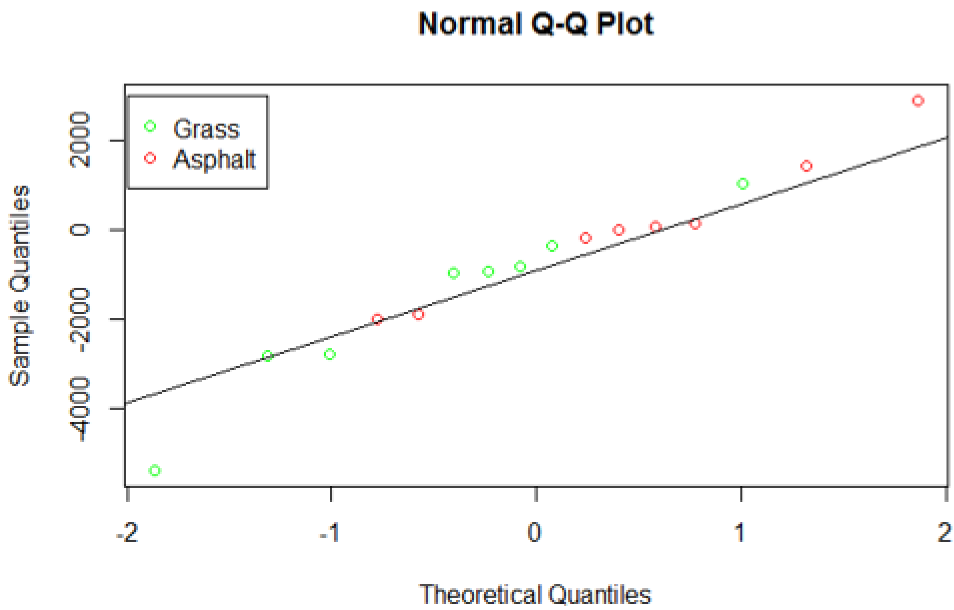

Fig. 9 **Normal Q–Q plot of differences in energy expended for each configuration. The line shows the path a normal distribution would follow.**

In Section 4.1, we conducted tests to determine whether we should expect a difference in mean energy use between the 2 planners. We conducted these tests using all pairs of runs, excluding the extreme values represented by configurations 5 and 14. The exclusion of these values does not mean they are in error and a review of the video suggests no measurement error in them; they simply represent extreme failures of the energy-efficient/minimum-distance planners in configurations 5 and 14. Such occasional failures are common in robotics and could represent failures by other aspects of the system—for example, perception or localization.

While the results with and without extreme values differ in their p-values, the confidence intervals are similar and both suggest that on average, energy-efficient planning really is more energy efficient as compared with the minimum-distance planner used for these tests.

5. Conclusions

The energy-efficient planning proposed in Ordonez et al.[1] and Gupta et al.[2] showed promise—demonstrating superior performance on grass, navigating with fewer collisions, and expending less energy on average. There was not a clear difference on asphalt, however. Overall, there is evidence to suggest that on average, energy-efficient planning will save energy—at least by comparison with the version of minimum-distance planning with which it was compared here. It would be beneficial to examine the differences in performance when the platforms operate on slopes, where energy-efficient planning could be expected to have substantially greater performance. The test courses were small but it is reasonable to expect that the advantages in energy expenditure may hold up on a larger scale, provided a map is available for the energy-efficient planner to use. On a larger scale, however, the assumption that terrain can be approximated by a plane would not be reasonable and an extension of the algorithm will be required for this situation. We would expect that on a more cluttered course, with a good map, the energy-efficient planner would have a greater advantage, while a course that is more open, or one without an available map, would reduce or eliminate the advantages of energy-efficient planning.

6. References

1. Ordonez C, Gupta N, Reese B, Seegmiller N, Kelly A, Collins E. Learning of skid-steered kinematic and dynamic models for motion planning. Robotics and Autonomous Sys. 2015.

2. Gupta N, Ordonez C, Collins E. Dynamically feasible, energy efficient motion planning for skid-steered vehicles. Autonomous Robots. 2016:1–19.

3. Dean R. Common world model for unmanned systems. Proc. SPIE; c2013.

4. Lennon C, Bodt B, Childers M, Dean R, Oh J, DiBerardino C, Keegan T. RCTA capstone assessment. Proc. SPIE; c2015.

Appendix. Data from Runs

Table A-1 Recorded data from all runs

Terrain	Block	Planner	Tracking re-plan	Collision re-plan	Actual energy	Collision	Run	Time
Asphalt	1	Distance	5	4	6,838.24	No	1	20
Asphalt	1	Energy	4	9	6,926.58	No	2	22
Asphalt	2	Energy	3	6	5,972.92	No	3	30
Asphalt	2	Distance	3	14	7,849.53	No	4	26
Asphalt	3	Energy	6	8	6,651.2	No	5	25
Asphalt	3	Distance	5	18	8,625.42	No	6	26
Asphalt	4	Energy	4	5	7,058.7	No	7	27
Asphalt	4	Distance	2	25	7,032.11	No	8	26
Asphalt	5	Distance	3	4	4,305.44	No	9	26
Asphalt	5	Energy	3	5	7,224.85	No	10	25
Asphalt	6	Distance	5	1	5,211.67	No	11	23
Asphalt	6	Energy*	8	0	4,130.46	No	12	26
Asphalt	6	Energy	5	4	5,044.96	No	13	24
Asphalt	7	Energy*	9	3	3,491.94	No	14	28
Asphalt	7	Distance	3	8	5,844.71	No	15	28
Asphalt	7	Energy	5	0	5,983.67	No	16	24
Asphalt	8	Distance	3	10	5,603.29	No	17	24
Asphalt	8	Energy	3	13	7,056.72	No	18	24
Grass	9	Energy	4	6	4,695.81	No	19	24
Grass	9	Distance	3	14	5,667	Yes	20	NA
Grass	10	Distance	7	1	5,101.76	No	21	23
Grass	10	Energy	5	2	4,272.85	No	22	24
Grass	11	Distance	7	15	6,631.67	Yes	23	26
Grass	11	Energy	6	2	3,856.34	No	24	23
Grass	12	Distance	4	3	3,512.46	No	25	22
Grass	12	Energy	2	6	2,589.58	No	26	21
Grass	13	Distance	5	18	7,472.84	Yes	27	25
Grass	13	Energy	2	15	4,653.34	No	28	22
Grass	14	Energy	4	3	3,530.1	No	29	24
Grass	14	Distance	6	13	8,924.45	Yes	30	28
Grass	15	Energy	6	8	5,079.44	No	31	25
Grass	15	Energy*	6	17	4,622.06	Yes	32	26
Grass	15	Distance	3	7	4,045.69	Yes	33	25
Grass	16	Energy*	6	10	4,651.31	No	34	25
Grass	16	Distance	4	6	6,148.35	No	35	27
Grass	16	Energy	3	11	5,804.02	No	36	29

Table A-2 Definitions for all recorded variables

Variable	Definition
Terrain	Type of terrain (grass vs. asphalt)
Block	The configuration of barrels
Planner	The type of planner used (minimum distance, and energy efficient with or without learning)
Tracking re-plan	The number of times re-planning was done because the trajectory was different from that originally planned
Collision re-plan	The number of times re-planning was done because the trajectory was different from that originally planned
Actual energy	Joules used during the run
HitBarrel	Denotes a collision with a barrel (1) or no collisions (0)
Run	The number of the run
Time	The length of time it took the run to complete

INTENTIONALLY LEFT BLANK.

List of Symbols, Abbreviations, and Acronyms

ANOVA	analysis of variance
ARL	US Army Research Laboratory
EE	energy efficient
EE Default	energy efficient without terrain learning
MD	minimum distance
Q–Q	quantile–quantile
RCTA	Robotics Collaborative Technology Alliance

1 (PDF)	DEFENSE TECHNICAL INFORMATION CTR DTIC OCA
2 (PDF)	DIRECTOR US ARMY RSRCH LAB RDRL CIO L IMAL HRA MAIL & RECORDS MGMT
1 (PDF)	GOVT PRINTG OFC A MALHOTRA
2 (PDF)	DIR USARL RDRL VTA C LENNON M CHILDERS